

LOAD BALANCING STORAGE SYSTEM

The present application is a continuation of application Serial No. 09/836,458, filed April 18, 2001, the contents of which are incorporated herein by reference.

BACKGROUND OF THE INVENTION

[0001] The present invention relates to a storage system operating in communication with a host unit; and, more particularly, the invention relates to the connecting configuration of recording media having the structure of a loop connecting plural recording media.

[0002] Recently, with the increased interest in the concept of a Storage Area Network (SAN), the loop connection using an optical fiber channel is rapidly being adopted as an interface to connect a host unit with various devices. The structure of the connections between components inside a disk array unit, which does not use a bus connection, but uses a loop connection having a fiber optic channel, is becoming popular (for example, a loop connection between a controller and recording media using a fiber optic channel). For instance, Japanese patent publication No. Kokai 11-338646 discloses a disk array unit which does not use a bus, but the switching network structure for the connection between components.

SUMMARY OF THE INVENTION

[0003] The above-described patent publication is directed to the connection between the components with a single network. In a storage system connecting

components with plural loops, such as this network, the following considerations will be important for mounting and setting up the recording media.

1. In case the storage system is a disk array system, the data from the host unit is distributed, generated and relocated to the data by the Redundant Array of Inexpensive (Independent) Disks (RAID) method. The data from the host unit is also recorded dispersedly into plural recording media which are set as a RAID group. Therefore, it is important to optimize the access paths and the load balancing to each recording medium. 2. There are plural loops of the recording media. Therefore, it is important to optimize the access path and the load balancing to these loops.

[0004] However, the above-described patent publication discloses nothing about the optimization of the access paths and the load balancing in connecting the recording media with multipath loops for the recording media.

[0005] As a method of load balancing optimally, it is desirable to mount the recording media at an optimum location, considering the loop for the recording medium, at the time of mounting the recording medium. In addition, at the time of forming a RAID group with the mounted recording media, it is desirable to select the recording media connected to the optimum loop for the recording media in order to build the RAID group. However, the selection of the mounting location or the selection of the recording media while taking into consideration the problem of load balancing is difficult for a person other than those who are familiar with the internal structure of the controller, such as the connecting location of the loop. Even when an operator who is in charge of the mounting and the setting executes such operation, there is a high possibility that such person will fail in distributing the recording media with a uniform load.

[0006] The object of the present invention is to provide a recording media mounting system that enables the adequate setting of access paths and load balancing without specifically considering the mounting location.

[0007] As a configuration to attain the above-described object, the storage system of the present invention consists of two or more recording media storing data received from a host unit, a controller controlling the data transfer between the recording media and the host unit, and plural loops, located between the controller and the recording media, connecting the different recording media. A recording medium is connected to an adjacent different loop than that of another recording medium, and the number of the recording media connected to each loop can be equalized by simply connecting the recording media to successive loops in turn.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] FIG. 1 is a block diagram of the total structure of the present invention.

[0009] FIG. 2 is a block diagram of the configuration inside a disk controller.

[0010] FIG. 3 is a schematic diagram illustrating an example of the connection of hard disk drives with-loops.

[0011] FIG. 4 is a schematic diagram illustrating an example of the connection of hard disk drives with loops.

[0012] FIG. 5 is a schematic diagram illustrating an example of the connection of hard disk drives with loops.

[0013] FIG. 6 is a schematic diagram illustrating an example of the connection of the loops with a multiplicity of disk drives.

[0014] FIG. 7 is a perspective view of a disk storage subsystem with the front panel removed.

DETAILED DESCRIPTION OF THE INVENTION

[0015] A preferred embodiment of the present invention will be described with reference to the drawings.

[0016] FIG. 1 is an example of a configuration of the storage system, which uses hard disk- drives as the recording medium, in accordance with an embodiment of the present invention. Hereafter, in a disk array system where both the connections with a host unit and the connections with a group of recording media are optical fiber channel connections, a loop for performing communication and data transfer by connecting the controller with the group of recording media will be referred to as a back-end loop, and a loop for connecting with a host unit(s) will be referred to as a front-end loop for purposes of description.

[0017] In FIG. 1, a disk controller 112 is provided with a port 111, which is an interface to connect the disk controller with host units 100-1, 100-2 and 100-3, that operate as host units or central processing units, through a Storage Area Network (SAN), and ports 150 and 151, that serve as interfaces to connect the disk controller 112 to back-end loops 140 and 141, respectively, through which the disk controller is connected with the hard disk drives 120 that operate as the recording media.

[0018] The disk controller 112 controls the transfer of data written from the host units 100 to the hard disk drives 120, or the transfer of data read from the

hard disk drives 120 to the host units 100. In a word, the disk controller 112 controls the whole disk storage subsystem 110.

[0019] This disk storage subsystem 110 consists of the hard disk drives 120-0 through 120-9; connecting bays 130-0 through 130-9, which function as connectors connecting the hard disk drives detachably with the back-end loops that are formed by the optical fiber channels; the back-end loops 140 and 141, which connect the hard disk drives to the disk controller; and the disk controller 112 itself for controlling the hard disk drives. The back-end loops- 140 and 141 are electrically connected through the connecting bays 130-0 through 130-9 with the hard disk drives 120-0 through 120-9. It is possible to connect the hard disk drives 120-0 through 120-9 to the connecting bays 130-0 through 130-9 independently, and it is not necessary to mount connecting drives in all of the connecting bays.

[0020] The disk storage subsystem 110 uses the loop interfaces of the optical fiber channels for all connections of the back-end loops 140 and 141, the connecting bays 130-0 through 130-9, the hard disk drives 120-0 through 120-9, and the interfaces 150 and -151. It is needless to say that the back-end loops 140 and 141, that operate as communication paths, are not limited to optical fiber channels.

[0021] The back-end loop is illustrated in the figure as an ellipse and with the respective lines from the ellipse extending to respective ones of the connecting bays 130. In the figure, though each connection is shown by a line for ease of description, the connection from the ellipse to the connecting bay 130 also forms a loop which is electrically detachable with respect to the loop indicated by the

ellipse. In the following description, each of the loops is shown by one line, except for the loops in FIG. 3. FIG. 3 shows the loop as a pair of lines.

[0022] Here, the port 150 is connected with the back-end loop 140, and the port 151 is connected with the back-end loop 141. The connecting bays 130-0, 130-2, 130-4, 130-6 and 130-8 are connected with the back-end loop 140, and the connecting bays 130-1, 130-3, 130-5, 130-7 and 130-9 are connected with the back-end loop 141. The mounting position of each connecting bay alternately connected with each back-end loop is, for example, mounted from the left side sequentially in the order of the consecutive connecting bay number, as shown in the figure. Therefore, in case of mounting the hard disk drives 120 to each connecting bay, regardless of the number of hard disk drives, it is easy to alternately allocate the hard disk drives 120 to the back-end loops 140 and 141 by mounting the hard disk drives 120 in the ascending order from the location of the connecting bay 130-0 or by mounting the hard disk drives 120 in the descending order from the location of the connecting bay 130-9. Therefore, the number of hard disk drives connected to each back-end loop can be shared equally. Connecting the hard disk drives 120 to each of the back-end loops 140 and 141 alternately is a method which results in optimal distribution of the load to the throughput of the back-end loop.

[0023] FIG. 2 shows the outline of the controller 112 of the disk storage subsystem 110. The controller 112 comprises a port 111 connected in communication with a host unit 100 to control protocol processing and data transfer with the host unit 100; ports 150 and 151 connected to the hard disk drive group 120 to control protocol processing and data transfer with the hard disk drive 120; a cache memory 202 for storing the transferred data temporarily;

the CPU 203 that operate as a processor for controlling each of these elements and the entire disk controller 112; and a bus 201 which provides data transfer between the ports and the cache memory and command transfer between the ports and the CPU.

[0024] When the CPU 203 of the disk storage subsystem 110 receives a data write instruction from the host unit 100, the CPU 203 temporarily stores the data sent from the host unit 100 via the port 111 in the cache memory 202 through the bus 201. The CPU 203 reports the end of the write operation to the host unit when the CPU 203 completes the storage of the data in the cache memory 202. Then, the processor in the port 150 or 151 writes the data stored in the cache memory 202 into the hard disk drive group 120 while being controlled with RAID technology.

[0025] When the CPU receives the data read instruction from the host unit 100, the CPU refers to the information managing the cache memory and decides whether the instructed data is stored in the cache memory 202. If the instructed data is stored in the cache memory 202, the CPU transfers the data to the host unit 100 through the port 111. If the data is not stored in the cache memory 202, the CPU copies the data from the hard disk drives 120 to the cache memory 202, and then transfers the data from the cache memory 202 to the host unit 100 via the port 111.

[0026] For ease of description, an example of a single host unit 100 is shown; however, it is possible to connect the port 111 to the SAN as shown in FIG. 1, and it is also possible to connect the ports 111 to two or more host units provided with two or more ports 111. The disk controller 112 can be not only bus

connected, but it can also connect the cache memory to each interface through an independent path.

[0027] Next, with reference to FIG. 3, a system comprising two RAID groups 160 and 161 will be explained, one of which is assigned four hard disk drives and the other of which is assigned two hard disk drives.

[0028] First, the mounting method will be explained. In a disk storage subsystem 110 shown in FIG. 3, the connecting bays 130-0, 130-3, 130-4, 130-7 and 130-8 are connected to the back-end loop 140, and the connecting bays 130-1, 130-2, 130-5, 130-6 and 130-9 are connected to the back-end loop 141. Forming pairs of hard disk drives 120, except for other parts, can decrease the number of loops from the back-end loop to the connecting bay.

[0029] With such a configuration of the loops and the connecting bays, mounting the hard disk drives 120 from the connecting bay 130-0 in the ascending order or from the connecting bay 130-9 in the descending order allows the fiber channel loops 140 and 141 to have the same number of the hard disk drives 120, when the total number is even, and to have the numbers of the hard disk drives 120 different by one drive, when the total number is odd. The load is balanced among the back-end loops. It is needless to say that the configuration of pairs of the hard disk drives 120, including the other parts, can provide an effectively balanced load.

[0030] Similarly to FIG. 1, the load can be balanced with the connection of the connecting bays 130-0, 130-2, 130-4, 130-6 and 130-8 to the back-end loop 140 and the connection of the connecting bays 130-1, 130-3, 130-5, 130-7 and 130-9 to the back-end loop 141.

[0031] Next, the configuration of the RAID groups 160 and 161, will be described. Four hard disk drives required by the RAID group 160 are allocated from the connecting bay 130-0 in the ascending order. Here, the hard disk drives are allocated alternately to the back-end loops 140 and 141, respectively. Similarly, a pair of required hard disk drives are allocated from the connecting bay 130-4 in the ascending order to the RAID group 161, and then the load is balanced among the back-end loops.

[0032] In this preferred embodiment of the present invention, a simple operation of allocating the required number of the hard disk drives from the connecting bay 130-0 in the ascending order can realize a mounting of the hard disk drives and a load balancing among the back-end loops of the hard disk drives inside the RAID group.

[0033] Next, the preferred embodiment of the expansion of the back-end loops, the connecting bays, and the hard disk drives will be explained with reference to FIG. 4. The components of the disk storage subsystem 110 are the components used in FIG. 1 and the disk storage subsystem 110 is connected with an extension unit 1110 containing back-end loops 142 and 143, connecting bays 131-0 through 131-9, and hard disk drives 121-0 through 121-9. The extension unit 1110 is connected with the disk storage subsystem 110 through the connection of the back-end loop 142 with the back-end loop 140 and the back-end loop 143 with the back-end loop 141. For the ease of description, the figure shows the connection between the back-end loop in the disk storage subsystem 110 and the back-end loop in the extension unit 1110 as the connection of a single line; however, another loop is formed with the back-end

loop in the disk storage subsystem 110 and the back-end loop in the extension unit 1110.

[0034] In this case, the back-end loop and the connecting bay are connected as follows: the connecting bays 131-0, 131-2, 131-4, 131-6 and 131-8 connect to the back-end loop 142, and the connecting bays 131-1, 131-3, 131-5, 131-7 and 131-9 connect to the back-end loop 143. In this case, the above-described connection makes it possible to count or set up the hard disk drives 120 to each connecting bay.

[0035] The following describes an example of a configuration of two RAID groups. One RAID group is allocated with twelve of the hard disk drives 120 and the other RAID group is allocated with four of the hard disk drives 120.

[0036] The RAID group is mounted with the necessary twelve hard disk drives from the connection bay 130-0 in the ascending order. In a case where the hard disk drives 120 cannot be connected due to the failure of the connecting bay 130-0 or some other reason, the mounting of the hard disk drives 120 does not have to start from the connecting bay 130-0, but it can start from any other of the connecting bays 130 sequentially. The following describes the configuration of the RAID groups 162 and 163. The twelve hard disk drives 120 used for the RAID group 162 are sequentially allocated from the connecting bay 130-0 or any mounted connecting bay.

[0037] The sequential allocation from the connecting bays allocates the hard disk drives 120 alternately to the connected group of the back-end loops 140 and 142 and the connected group of the back-end loops 141 and 143; therefore, the load can be balanced between the back-end loops. As in the case of the RAID group 162, the RAID group 163 can be allocated with the four necessary hard

disk drives 120 from the connecting bay 131-2 sequentially to balance the load between the back-end loops. In this preferred embodiment of the present invention, a simple rule of sequentially allocating the necessary number of hard disk drives 120 from the connecting bay 130 can realize a mounting of a hard disk drives 120 and the load balancing among the loops of the hard disk drives inside the RAID group.

[0038] FIG. 5 shows the preferred embodiment of the present invention for a disk storage subsystem provided with three back-end loops. The components of the disk storage subsystem 110 are similar to the components of the disk storage subsystem used in FIG. 1. The main difference from the preferred embodiment of FIG. 1 is that there are three systems in the connection between ports and back-end loops, and between back-end loops and connecting bays. Each of the ports 150, 151 and 152 is connected with the back-end loops 140, 141 and 144, respectively, the back-end loop 140 is connected with the connecting bays 130-0, 130-3, 130-6 and 130-9, the back-end loop 141 is connected with the connecting bays 130-1, 130-4 and 130-7 and the back-end loop 144 is connected with the connecting bays 130-2, 130-5 and 130-8.

[0039] Mounting the hard disk drives 120 to each of the connecting bays 130 connects the adjacent hard disk drives 120 to a different back-end loop and port. In this preferred embodiment, the theory of the load balancing is similar to the case of the two back-end loops. Three back-end loops will reduce the amount of the load per one back-end loop as compared with two back-end loops.

[0040] FIG. 6 shows an example of a storage system provided with plural controllers and loops. The data is generally transferred between the back-end loop 140 or 141 and each of the connecting bays 130 through the optical fiber

channel shown with a solid line in FIG. 6. That is, the data transfer under normal condition is as in the case of FIG. 1. In a case where a failure has occurred in the optical fiber channel between the back-end loop and each of the connecting bays 130, the data is transferred through the optical fiber channel shown with a dotted line using a different back-end loop from the normal one.

[0041] To be more specific, the normal connection route for the hard disk drive 120-1 is from the host unit 100 through the port III and the port 151 and via the back-end loop 141 to the connecting bay 130-1; however, in case of a failure between the back-end loop 141 and the connecting bay 130-1, the connection route for the hard disk drive 120-1 will be changed to another route from the host unit 100 through the port 111, the port 150, and the back-end loop 140 to the connecting bay 130-1, so that the data can be transferred without giving the user any reduction of the failure.

[0042] In case of a failure in one port or in one back-end loop, the other port and the other back-end loop make data transfer possible between the host unit 100 and the hard disk drive 120.

[0043] Though the above-described preferred embodiments of the present invention have been explained with reference to a hard disk drive as a recording medium, the recording medium can be an optical disk or a magneto-optical disk, or can be either of a tape unit or a card type memory, such as a flash memory.

[0044] The above-described disk storage subsystem 110 has, for instance, the external appearance as shown in FIG. 7. The disk storage subsystem 110 has a rectangular frame 301, a front panel 302 with ventilating holes, which is located in front of the disk storage subsystem 110 and covers the front surface, and the connecting bays 130-0 through 130-9 which are arranged in numerical

order to mount the hard disk drives 120 on the upper side of the front panel 302. Mounting into these connecting bays 130 allows the hard disk drives to be located in a row consisting of ten independent drives in sequence from 120-0 through 120-9.

[0045] The extension unit 1110 explained with reference to FIG. 4 has a shape for accommodating the connecting bays 131 for mounting the hard disk drives 121 on the upper side of the frame 301, and the back-end loop on the disk array unit 110 can be connected with the back-end loop on the extension unit 1110 on the rear panel shown in FIG. 7.

[0046] Therefore, the hard disk drives can be easily expanded by stacking the extension unit 1110 and the disk storage subsystem 110. In stacking the units, it is possible to mount the disk storage subsystem 110 and the extension unit 1110 into a rack, or to cover the disk storage subsystem 110 and the extension units 1110 with a frame having a size according to the number of the hard disk drives the user desires.

[0047] According to the present invention, as described above, in the disk storage subsystem provided with two or more back-end loops for connecting the hard disk drives, the load balancing in each loop and the setting work for mounting can be facilitated by connecting the back-end loops to the connecting bays which mount the hard disk drives with an adequate pattern so that the connecting bays conform to the mounting order of the hard disk drives.

[0048]